

New Developments in FRE: dual runs and libraries

FMS Developers' Forum

V. Balaji¹ A. Langenhorst² N. Zadeh²

¹Princeton University

²RSIS/GFDL

28 March 2007

- 1 New Developments in FRE
 - Dual run capability in FRE
 - FRE Redesign: the new `fremake`
 - FRE Redesign: modular site configuration
 - FRE Redesign: next steps
- 2 Nalanda: the April patch

1 New Developments in FRE

- Dual run capability in FRE
- FRE Redesign: the new `fremake`
- FRE Redesign: modular site configuration
- FRE Redesign: next steps

2 Nalanda: the April patch

1 New Developments in FRE

- Dual run capability in FRE
- FRE Redesign: the new `fremake`
- FRE Redesign: modular site configuration
- FRE Redesign: next steps

2 Nalanda: the April patch

Dual-running an experiment in FRE...

Recent concerns about the integrity of our simulations have led to the request to add the capability to rerun existing simulations exactly. This is what is now known as the *dual run*.

- There are no changes needed to the FRE schema ("XML file").
`frerun -u` of any existing run will rerun any past or current run exactly as specified in the FRE, but without clobbering any output files.
- The second run does not incur charges against the group allocation. Also, no post-processing is performed.

Dual-running an experiment in FRE...

Recent concerns about the integrity of our simulations have led to the request to add the capability to rerun existing simulations exactly. This is what is now known as the *dual run*.

- There are no changes needed to the FRE schema ("XML file").
`frerun -u` of any existing run will rerun any past or current run exactly as specified in the FRE, but without clobbering any output files.
- The second run does not incur charges against the group allocation. Also, no post-processing is performed.

Dual-running an experiment in FRE... (contd.)

Recent concerns about the integrity of our simulations have led to the request to add the capability to rerun existing simulations exactly. This is what is now known as the *dual run*.

- Use `frecheck` to perform integrity checks (bitwise comparisons or restart and raw history files) between the run and the dual run.
not working yet for dual runs...
- In future, the FRE schema will evolve toward explicit support of multiple "realizations" of a run. The FRE DB will have the capability to return the exact difference in configuration between two realizations.

More info: <http://cobweb.gfdl.noaa.gov/~vb/fre>

Dual-running an experiment in FRE... (contd.)

Recent concerns about the integrity of our simulations have led to the request to add the capability to rerun existing simulations exactly. This is what is now known as the *dual run*.

- Use `frecheck` to perform integrity checks (bitwise comparisons or restart and raw history files) between the run and the dual run.
not working yet for dual runs...
- In future, the FRE schema will evolve toward explicit support of multiple "realizations" of a run. The FRE DB will have the capability to return the exact difference in configuration between two realizations.

More info: <http://cobweb.gfdl.noaa.gov/~vb/fre>

Amy's dual-run FAQ...

- How do I start a dual run for a new experiment or as a reproducibility test for an old experiment?
 - `frerun`, then `frerun -u`. The second run appears in the directory `.../experiment/1`.
- How do I rerun just a subset of a previous experiment? Will I be charged for the hours it takes to run?
 - The run will start from the specified `<initCond>` file and run until you stop it.
 - Any job submitted with `qsub -l repro -A repro` is considered to be a dual-run job and will not incur charges.

More info: <http://cobweb.gfdl.noaa.gov/~vb/fre>

Amy's dual-run FAQ...

- How do I start a dual run for a new experiment or as a reproducibility test for an old experiment?
 - `frerun`, then `frerun -u`. The second run appears in the directory `.../experiment/1`.
- How do I rerun just a subset of a previous experiment? Will I be charged for the hours it takes to run?
 - The run will start from the specified `<initCond>` file and run until you stop it.
 - Any job submitted with `qsub -l repro -A repro` is considered to be a dual-run job and will not incur charges.

More info: <http://cobweb.gfdl.noaa.gov/~vb/fre>

Amy's dual-run FAQ...

- How do I start a dual run for a new experiment or as a reproducibility test for an old experiment?
 - `frerun`, then `frerun -u`. The second run appears in the directory `.../experiment/1`.
- How do I rerun just a subset of a previous experiment? Will I be charged for the hours it takes to run?
 - The run will start from the specified `<initCond>` file and run until you stop it.
 - Any job submitted with `qsub -l repro -A repro` is considered to be a dual-run job and will not incur charges.

More info: <http://cobweb.gfdl.noaa.gov/~vb/fre>

Amy's dual-run FAQ...

- How do I start a dual run for a new experiment or as a reproducibility test for an old experiment?
 - `frerun`, then `frerun -u`. The second run appears in the directory `.../experiment/1`.
- How do I rerun just a subset of a previous experiment? Will I be charged for the hours it takes to run?
 - The run will start from the specified `<initCond>` file and run until you stop it.
 - Any job submitted with `qsub -l repro -A repro` is considered to be a dual-run job and will not incur charges.

More info: <http://cobweb.gfdl.noaa.gov/~vb/fre>

Amy's dual-run FAQ...

- How do I start a dual run for a new experiment or as a reproducibility test for an old experiment?
 - `frerun`, then `frerun -u`. The second run appears in the directory `.../experiment/1`.
- How do I rerun just a subset of a previous experiment? Will I be charged for the hours it takes to run?
 - The run will start from the specified `<initCond>` file and run until you stop it.
 - Any job submitted with `qsub -l repro -A repro` is considered to be a dual-run job and will not incur charges.

More info: <http://cobweb.gfdl.noaa.gov/~vb/fre>

Amy's dual-run FAQ... (contd.)

- How do I tell what hosts and cpusets my previous experiments ran on?
 - The job epilogue in the `stdout` file will report the hostname and the *logical* cpuset. If the script does `setenv MPI_DSM_VERBOSE` you get a list of brick identifiers which (I believe) are hardware IDs.
 - You can assign the job to a specific host with `qsub -l ic7`, for instance. (Type `qconf -scl` for a list).

More info: <http://cobweb.gfdl.noaa.gov/~vb/fre>

Amy's dual-run FAQ... (contd.)

- How do I tell what hosts and cpusets my previous experiments ran on?
 - The job epilogue in the `stdout` file will report the hostname and the *logical* cpuset. If the script does `setenv MPI_DSM_VERBOSE` you get a list of brick identifiers which (I believe) are hardware IDs.
 - You can assign the job to a specific host with `qsub -l ic7`, for instance. (Type `qconf -scl` for a list).

More info: <http://cobweb.gfdl.noaa.gov/~vb/fre>

Amy's dual-run FAQ... (contd.)

- How do I tell what hosts and cpusets my previous experiments ran on?
 - The job epilogue in the `stdout` file will report the hostname and the *logical* cpuset. If the script does `setenv MPI_DSM_VERBOSE` you get a list of brick identifiers which (I believe) are hardware IDs.
 - You can assign the job to a specific host with `qsub -l ic7`, for instance. (Type `qconf -scl` for a list).

More info: <http://cobweb.gfdl.noaa.gov/~vb/fre>

Amy's dual-run FAQ... (contd.)

- How do I compare the results of the two runs?
 - `frecheck` does a comprehensive check; **not working yet for dual runs...**
 - `resdiff` compares restart files using `cmp`;
 - `hisdifff` compares history files using `nccmp` (slower but cleverer than `cmp`).
- For which experiments should I perform dual runs?
 - It's up to you! It's not unreasonable to considering rerunning any production run. If you consider `i1c6` to be under a bigger cloud of suspicion, you can ask Ops for list of jobs that ran on it, and even on specific `cpusets` within it.

More info: <http://cobweb.gfdl.noaa.gov/~vb/fre>

Amy's dual-run FAQ... (contd.)

- How do I compare the results of the two runs?
 - `frecheck` does a comprehensive check; **not working yet for dual runs...**
 - `resdiff` compares restart files using `cmp`;
 - `hisdifff` compares history files using `nccmp` (slower but cleverer than `cmp`).
- For which experiments should I perform dual runs?
 - It's up to you! It's not unreasonable to considering rerunning any production run. If you consider `i26` to be under a bigger cloud of suspicion, you can ask Ops for list of jobs that ran on it, and even on specific `cpusets` within it.

More info: <http://cobweb.gfdl.noaa.gov/~vb/fre>

Amy's dual-run FAQ... (contd.)

- How do I compare the results of the two runs?
 - `frecheck` does a comprehensive check; **not working yet for dual runs...**
 - `resdiff` compares restart files using `cmp`;
 - `hstdiff` compares history files using `nccmp` (slower but cleverer than `cmp`).
- For which experiments should I perform dual runs?
 - It's up to you! It's not unreasonable to considering rerunning any production run. If you consider `i26` to be under a bigger cloud of suspicion, you can ask Ops for list of jobs that ran on it, and even on specific `cpusets` within it.

More info: <http://cobweb.gfdl.noaa.gov/~vb/fre>

Amy's dual-run FAQ... (contd.)

- How do I compare the results of the two runs?
 - `frecheck` does a comprehensive check; **not working yet for dual runs...**
 - `resdiff` compares restart files using `cmp`;
 - `hstdiff` compares history files using `nccmp` (slower but cleverer than `cmp`).
- For which experiments should I perform dual runs?
 - It's up to you! It's not unreasonable to considering rerunning any production run. If you consider `i.e6` to be under a bigger cloud of suspicion, you can ask Ops for list of jobs that ran on it, and even on specific `cpusets` within it.

More info: <http://cobweb.gfdl.noaa.gov/~vb/fre>

Amy's dual-run FAQ... (contd.)

- How do I compare the results of the two runs?
 - `frecheck` does a comprehensive check; **not working yet for dual runs...**
 - `resdiff` compares restart files using `cmp`;
 - `hstdiff` compares history files using `nccmp` (slower but cleverer than `cmp`).
- For which experiments should I perform dual runs?
 - It's up to you! It's not unreasonable to considering rerunning any production run. If you consider `ic6` to be under a bigger cloud of suspicion, you can ask Ops for list of jobs that ran on it, and even on specific `cpusets` within it.

More info: <http://cobweb.gfdl.noaa.gov/~vb/fre>

Amy's dual-run FAQ... (contd.)

- How do I compare the results of the two runs?
 - `frecheck` does a comprehensive check; **not working yet for dual runs...**
 - `resdiff` compares restart files using `cmp`;
 - `hstdiff` compares history files using `nccmp` (slower but cleverer than `cmp`).
- For which experiments should I perform dual runs?
 - It's up to you! It's not unreasonable to considering rerunning any production run. If you consider `ic6` to be under a bigger cloud of suspicion, you can ask Ops for list of jobs that ran on it, and even on specific cpusets within it.

More info: <http://cobweb.gfdl.noaa.gov/~vb/fre>

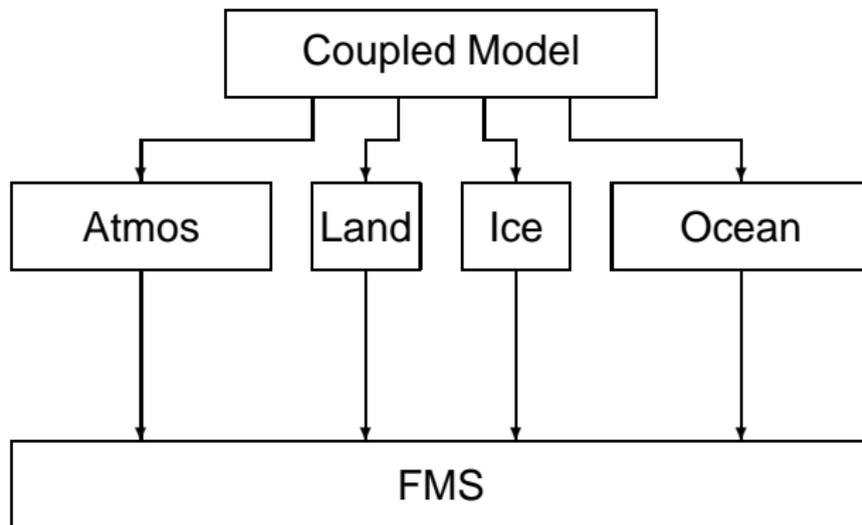
1 New Developments in FRE

- Dual run capability in FRE
- **FRE Redesign: the new `fremake`**
- FRE Redesign: modular site configuration
- FRE Redesign: next steps

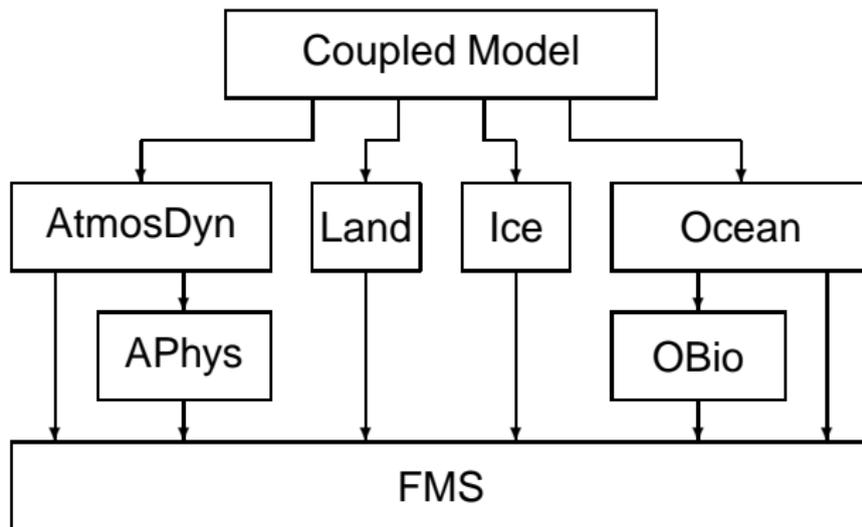
2 Nalanda: the April patch

FRE in components...

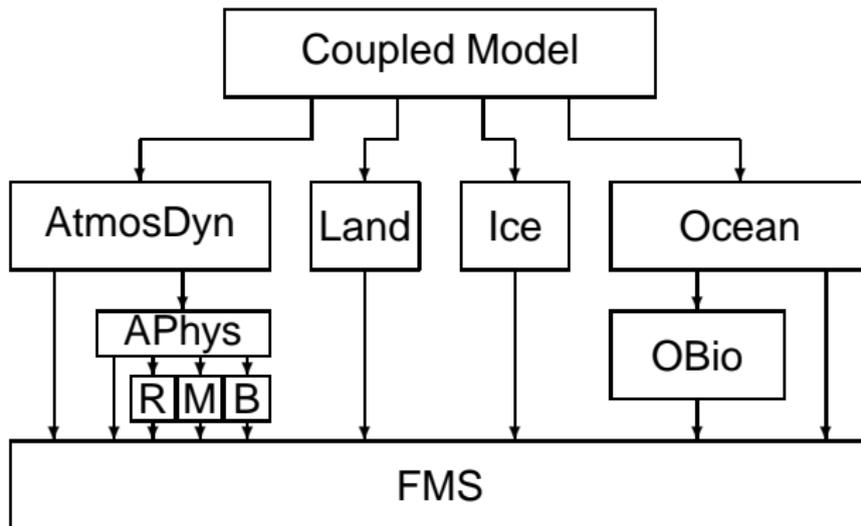
The modular structure of FMS is now reflected in the FRE schema.



Or...



Or even...



FRE Redesign: organization by component.

- As part of the FRE redesign, the schema will be entirely organized by component, which can be defined at any granularity, as shown. The advantage is that you can "fold away" and treat most model components as a "black box" during development.
- Pre-compiled libraries will be made available for all components, at various levels of granularity (beginning with `libfms.a` for the shared code).

More info: <http://cobweb.gfdl.noaa.gov/~vb/fre>

FRE Redesign: organization by component.

- As part of the FRE redesign, the schema will be entirely organized by component, which can be defined at any granularity, as shown. The advantage is that you can "fold away" and treat most model components as a "black box" during development.
- Pre-compiled libraries will be made available for all components, at various levels of granularity (beginning with `libfms.a` for the shared code).

More info: <http://cobweb.gfdl.noaa.gov/~vb/fre>

The new fremake

Features: **libraries** and **multiple compiles from the same source**.

Component syntax:

```
<component name="fms" paths="shared">
  <source versionControl="cvs" root="/home/fms/cvs">
    <codeBase version="latest">shared</codeBase>
  </source>
  <compile>
    <cppDefs>-Duse_libMPI -Duse_netCDF</cppDefs>
  </compile>
  <compile target="debug">
    <cppDefs>-Duse_libMPI -Duse_netCDF</cppDefs>
    <mkmfTemplate>mkmf.debug.ia64</mkmfTemplate>
  </compile>
</component>
```

(1)

More info: <http://cobweb.gfdl.noaa.gov/~vb/fre>

Dependencies

```
<component name="aphys" requires="fms">  
  .  
  .  
  .  
</component>  
<component name="adyn" requires="fms aphys">  
  .  
  .  
  .  
</component>
```

(2)

More info: <http://cobweb.gfdl.noaa.gov/~vb/fre>

Pre-compiled libraries

```
<component name="fms" paths="shared">
  <library path="../../../m45_am2p14/exec/libfms.a"
    headerDir="../../../m45_am2p14/exec" />
  <source versionControl="cvs" root="/home/fms/cvs">
    <codeBase version="latest">shared</codeBase>
  </source>
  <compile>
    <cppDefs>-Duse_libMPI -Duse_netCDF</cppDefs>
  </compile>
</component>
```

(3)

We will make available libraries for three modes, *production*, *regression-testing* and *debugging*.

More info: <http://cobweb.gfdl.noaa.gov/~vb/fre>

1 New Developments in FRE

- Dual run capability in FRE
- FRE Redesign: the new `fremake`
- **FRE Redesign: modular site configuration**
- FRE Redesign: next steps

2 Nalanda: the April patch

FRE site configuration

- As part of the redesign, the FRE scripts themselves will contain no hard-coded paths or site-specific information.
- You will be able to checkout and modify the site configuration, e.g by adding your own make templates.

```
fre_setup $HOME/fre  
will create, under $HOME/fre, directories bin/, lib/, site/,  
xml/, data/, ...
```

- `source $HOME/fre/site/fre.cshrc` will create your FRE environment.

More info: <http://cobweb.gfdl.noaa.gov/~vb/fre>

FRE site configuration

- As part of the redesign, the FRE scripts themselves will contain no hard-coded paths or site-specific information.
- You will be able to checkout and modify the site configuration, e.g by adding your own make templates.

```
fre_setup $HOME/fre
```

will create, under `$HOME/fre`, directories `bin/`, `lib/`, `site/`, `xml/`, `data/`, ...

- `source $HOME/fre/site/fre.cshrc` will create your FRE environment.

More info: <http://cobweb.gfdl.noaa.gov/~vb/fre>

FRE site configuration

- As part of the redesign, the FRE scripts themselves will contain no hard-coded paths or site-specific information.
- You will be able to checkout and modify the site configuration, e.g by adding your own make templates.

```
fre_setup $HOME/fre
```

will create, under `$HOME/fre`, directories `bin/`, `lib/`, `site/`, `xml/`, `data/`, ...

- `source $HOME/fre/site/fre.cshrc` will create your FRE environment.

More info: <http://cobweb.gfdl.noaa.gov/~vb/fre>

1 New Developments in FRE

- Dual run capability in FRE
- FRE Redesign: the new `fremake`
- FRE Redesign: modular site configuration
- FRE Redesign: next steps

2 Nalanda: the April patch

Next steps in FRE

- Full modularization:

```
<component>  
  <source>, <compile>, <input>, <pp>  
</component>  
...  
<component name="coupler">  
  <source>, <compile>, <input>, <run>, <pp>  
</component>
```

(4)

- Integration to database.
- file integrity using checksums.
- use of /ptmp
- triggers based on model time.

Next steps in FRE

- Full modularization:

```
<component>  
  <source>, <compile>, <input>, <pp>  
</component>  
...  
<component name="coupler">  
  <source>, <compile>, <input>, <run>, <pp>  
</component>
```

(4)

- Integration to database.

- file integrity using checksums.
- use of /ptmp
- triggers based on model time.

Next steps in FRE

- Full modularization:

```
<component>
  <source>, <compile>, <input>, <pp>
</component>
...
<component name="coupler">
  <source>, <compile>, <input>, <run>, <pp>
</component>
```

(4)

- Integration to database.
- file integrity using checksums.
- use of /ptmp
- triggers based on model time.

Next steps in FRE

- Full modularization:

```
<component>
  <source>, <compile>, <input>, <pp>
</component>
...
<component name="coupler">
  <source>, <compile>, <input>, <run>, <pp>
</component>
```

(4)

- Integration to database.
- file integrity using checksums.
- use of /ptmp
- triggers based on model time.

Next steps in FRE

- Full modularization:

```
<component>  
  <source>, <compile>, <input>, <pp>  
</component>  
...  
<component name="coupler">  
  <source>, <compile>, <input>, <run>, <pp>  
</component>
```

(4)

- Integration to database.
- file integrity using checksums.
- use of /ptmp
- triggers based on model time.

1 New Developments in FRE

- Dual run capability in FRE
- FRE Redesign: the new `fremake`
- FRE Redesign: modular site configuration
- FRE Redesign: next steps

2 Nalanda: the April patch

- Niki will describe some issues in the `nalanda` release, and a patch to resolve them.
- Most issues revolve around an incomplete implementation of the conservation code and the modified shortwave fluxes.
- We plan a rapid test and release of the April patch. When it's released, we strongly urge all users to migrate their codes to it.

- Niki will describe some issues in the `nalanda` release, and a patch to resolve them.
- Most issues revolve around an incomplete implementation of the conservation code and the modified shortwave fluxes.
- We plan a rapid test and release of the April patch. When it's released, we strongly urge all users to migrate their codes to it.

- Niki will describe some issues in the `nalanda` release, and a patch to resolve them.
- Most issues revolve around an incomplete implementation of the conservation code and the modified shortwave fluxes.
- We plan a rapid test and release of the April patch. When it's released, we strongly urge all users to migrate their codes to it.